

PATENT APPLICATION

Storage System, and Method for Controlling the Same

Inventors: **Masataka INNAN**
Citizenship: Japan

Akinobu SHIMADA
Citizenship: Japan

Hideo TABUCHI
Citizenship: Japan

Toshio NAKANO
Citizenship: Japan

Assignee: **Hitachi, Ltd.**
6, Kanda Surugadai 4-chome
Chiyoda-ku, Tokyo, Japan
Incorporation: Japan

Entity: Large

STORAGE SYSTEM, AND METHOD FOR CONTROLLING THE SAME

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority upon Japanese Patent
5 Application No. 2002-271595 filed on September 18, 2002 in Japan,
which is herein incorporated by reference.

BACKGROUND OF THE INVENTION

Field of the Invention

10 The present invention relates to a method for controlling
a storage system, a storage controller and a storage system.

Description of the Related Art

The environment surrounding information systems is rapidly
changing such as the progress of IT technology and shifting
15 to broadband. Along with this, measures against the rapid
increase of the amount of data handled are seriously required.
Therefore, the storage systems operating in data centers are
being subjected to prompt shifting to larger capacities and
higher performances of the storage controllers such as disk
20 array apparatuses.

To increase the storage area, a disk array apparatus can
additionally be provided for example. However, accordingly as
the number of the disk array apparatuses installed increases
the load of management on the host computer side is inevitably
25 increased, which may affect running of applications on the host
computer side. Furthermore, recently, so-called
"virtualization" is in progress that is a technology for managing
a huge storage area efficiently without making users aware of
the hardware configuration as far as possible. However, the
30 virtualization may cause an increase of the load of processing

on the host computer side.

In view of such circumstances, it is expected that the need for the technology for reducing the load of processing on the host computer side will become greater and greater. A technology is thus being sought that ensures larger capacities and higher performances of storage controllers without increasing the load of processing on the host computer side.

SUMMARY OF THE INVENTION

The present invention was conceived on the basis of the above background and one object thereof is to provide a method for controlling a storage system, a storage controller and a storage system.

A storage system according to one aspect of the present invention to achieve the above and other objects comprises a host computer; a first storage controller connected communicably to the host computer, for receiving a data frame transmitted from the host computer and executing data input to and data output from a first storage device in response to a data input/output request described in the data frame; and a second storage controller connected communicably to the first storage controller, for receiving a data frame transmitted from the host computer to the first storage controller and relayed by the first storage controller and executing data input to and data output from a second storage device in response to a data input/output request described in the data frame, wherein when the first storage controller receives the data frame transmitted from the host computer the first storage controller relays the data frame to the second storage controller in response to information described in the data frame.

According to the present invention there can be provided a method for controlling a storage system, a storage controller and a storage system.

5 Features and objects of the present invention other than the above will become clear by reading the description of the present specification with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

10 For a more complete understanding of the present invention and the advantages thereof, reference is now made to the following description taken in conjunction with the accompanying drawings wherein:

Fig. 1 illustrates the configuration of a storage system according to one embodiment of the present invention;

Fig. 2 illustrates the structure of a data frame on a Fiber Channel according to one embodiment of the present invention;

Fig. 3 illustrates the structure of a frame header according to one embodiment of the present invention;

20 Fig. 4 illustrates the structure of FCP_CMND according to one embodiment of the present invention;

Fig. 5 is a flow chart illustrating the operation of a storage system according to one embodiment of the present invention when a data frame is transmitted from a host computer to a first storage controller;

Fig. 6 shows an LUN management table according to one embodiment of the present invention;

Fig. 7 shows an access limit management table according to one embodiment of the present invention;

30 Fig. 8 is a flow chart illustrating an access limit function

according to one embodiment of the present invention;

Fig. 9 shows a priority management table according to one embodiment of the present invention;

Fig. 10 is a flow chart illustrating a priority control
5 function according to one embodiment of the present invention;

Fig. 11 shows a pair management table according to one embodiment of the present invention; and

Fig. 12 is a flow chart illustrating the processing of the storage system according to one embodiment of the present
10 invention when a data frame is transmitted in which a data writing request for a main logic volume is described.

DETAILED DESCRIPTION OF THE INVENTION

At least the following matters will be made clear by the
15 explanation in the present specification and the description of the accompanying drawings.

[Summary of the Disclosure]

A method for controlling a storage system, the storage system comprises a host computer; a first storage controller
20 connected communicably to the host computer, for receiving a data frame transmitted from the host computer and executing data input to and data output from a first storage device in response to a data input/output request described in the data frame; and a second storage controller connected communicably
25 to the first storage controller, for receiving a data frame transmitted from the host computer to the first storage controller and relayed by the first storage controller and executing data input to and data output from a second storage device in response to a data input/output request described
30 in the data frame, wherein when the first storage controller

receives the data frame transmitted from the host computer the first storage controller relays the data frame to the second storage controller in response to information described in the data frame.

5 In this method, the second storage controller is connected to the first storage controller such that the data frame transmitted from the host computer is received temporarily by the first storage controller and is transmitted by the first storage controller to the second storage controller in response
10 to information described in the data frame. In this configuration, the host computer transmits the data frame to the first storage controller indiscriminately regardless of whether the data frame that the host computer is going to transmit is directed to the first storage controller or to the second
15 storage controller. Therefore, the host computer has no need to communicate directly with the second storage controller when transmitting the data frame and the load of processing in the host computer can be reduced by the amount to be produced by the direct communication, compared to the scheme in which a
20 host computer executes direct communications with the second storage controller.

A function for relaying the data frame to the second storage controller in response to information described in the data frame when the data frame transmitted from the host computer
25 has been received can be realized by, for example, a program that is run on the first storage controller.

The data frame is, for example, a data frame conforming to Fiber Channel Protocol. The relaying is executed by a Fiber Channel switch included in the first storage controller.

30 The information is at least one of information for

specifying a Fiber Channel port of the source of the data frame, information for specifying a Fiber Channel port of the destination of the data frame and information for specifying a storage device. Among these, the information for specifying a Fiber Channel port of the source is, for example, Source_ID described later. The information for specifying a Fiber Channel port of the destination is, for example, Destination_ID described later. The information for specifying a storage device is, for example, FCP_LUN described later.

The storage device is, for example, a logic volume organized in zones on a storage area provided by a disk drive or a physical storage area provided by a disk drive.

In this method, the first storage controller functions rather as a network switch which either receives by itself the data frame transmitted from the host computer or relays the data frame to the second storage controller. Then, because the first storage controller functions as a network switch, an applied operation can be easily provided in which the information and the functions possessed by the first storage controller are combined with the function as a network switch. Therefore, functions such as, for example, access limit function, priority control, double management of data, described later, can be easily provided. Not limited to these functions, various functions can be easily realized which are obtained by combining the function as the network switch and the functions that the first storage controller originally has as a storage controller.

Furthermore, when the virtualization technology is applied to the storage system, a plurality of storage devices are managed virtually unitarily as one device and/or management/operation of functions such as remote copy and data copy management are

consolidated. By being provided with such functions, influence on the resource consumption and the processing rate on the host computer side becomes inevitable. However, by making available the provision of the functions formed by combining the functions that the first storage controller originally has as a storage controller and the function as that of a network switch by providing the first storage controller with the function as that of a network switch as described above, it is possible to provide the functions that have been originally included in the host computer, to the side of the first storage controller, to minimize the resource consumption and the load of processing on the host computer side and to provide a system supporting the virtualization.

Furthermore, when a general-purpose switch such as a Fiber Channel switch is used as the network switch, the first storage controller and the second storage controller can be easily connected even though those storage controllers are different models or are from different manufacturers. Therefore, the storage system can be easily configured and full use of resources such as existing or obsolete storage controllers can be progressed. In addition, in embodiments of (i) basic operations, (ii) access limit function, (iii) priority control and (iv) data copy management, full use of a plurality of storage controllers can be progressed by connecting one storage controller provided with these functions to a plurality of storage controllers not provided with these functions.

In the data frame, as the information, at least one of information for specifying a Fiber Channel port of the source of the data frame, information for specifying a Fiber Channel port of the destination of the data frame and information for

specifying a storage device, is described. The first storage controller stores access limit information indicating whether an access of a Fiber Channel port of the source to a Fiber Channel port or a storage device of the destination is permitted or not permitted. Therefore, it can be arranged that, when the first storage controller receives the data frame from the host computer, the first storage controller refers to the access limit information to check whether an access of a Fiber Channel port of the source of the data frame to a Fiber Channel port of the destination or a storage device of the destination is permitted and relays the data frame to the second storage controller only when the access is permitted.

In this manner, the first storage controller judges based on the access limit information whether the first storage controller should transmit the data frame to the second storage controller and, thus, the first storage controller executes replacing the role of the second storage controller the process relating to the access limit of the data frame directed to the second storage controller. Therefore, it is enough that the function for processing an access limit is implemented only on the first storage controller, and the load of maintenance is reduced by the amount to be created by taking the trouble of providing such a mechanism to the second storage controller. The function of the first storage controller relating to the access limit is realized by, for example, a program that is run on the first storage controller.

Since there is no need for the second storage controller to be provided with access limit function, a storage controller having lower performance than that of the first storage controller can be used as the second storage controller and,

therefore, the introduction cost and the operation cost can be reduced. Furthermore, old types of storage controllers having lower performance compared to the newly introduced storage controllers can be fully used as second storage controllers. Yet furthermore, it is enough that the access limit management table is stored only in the first storage controller and, therefore, reduction of the load of maintenance created by the unified management is progressed.

Arrangement may be as follows: in the data frame, as the information, information for specifying a Fiber Channel port of the source of the data frame and information for specifying a Fiber Channel port of the destination of the data frame are described; the first storage controller stores information indicating the priority of data input/output process corresponding to the data frame by relating to combinations of a Fiber Channel port of the source and a Fiber Channel port of the destination; when the first storage controller receives the data frame from the host computer, the apparatus executes data input to and data output from the storage device connected to a port of the destination of the data frame according to the information indicating the priority. The information indicating the priority is, for example, information indicating the timing for executing the data input to and data output from the storage device and, more specifically, a delay time which will be described later.

In this mechanism, it can be considered that the first storage controller executes the priority control function replacing the role of the second storage controller. That is, by only implementing the priority control function on the first storage controller, a priority control can also be executed

to the data frame directed to the second storage controller and the load of operation imposed on the second storage controller can be reduced by the amount to be created by providing such a mechanism.

5 The function of the first storage controller relating to the priority control is realized by, for example, a program that is run on the first storage controller.

Because there is no need for the second storage controller to be provided with priority control function a storage controller having lower processing performance than that of
10 the first storage controller can be used as the second storage controller and, therefore, the introduction cost and the operation cost can be reduced. Furthermore, old types of storage controllers having lower performance compared to the newly
15 introduced storage controllers can be fully used as second storage controllers. Yet furthermore, it is enough that the priority management table is stored only in the first storage controller, thus achieving a reduction of the load of maintenance by the unified management.

20 A method for controlling a storage system, the storage system comprising a host computer; a first storage controller connected communicably to the host computer, for receiving a data frame transmitted from the host computer and executing data input to and data output from a first storage device in
25 response to a data input/output request described in the data frame; and a second storage controller connected communicably to the first storage controller, for receiving a data frame transmitted from the host computer to the first storage controller and relayed by the first storage controller and
30 executing data input to and data output from a second storage

device in response to a data input/output request described in the data frame, wherein the first storage controller stores a copy of data stored in the first storage device which is the object of data input/output effected by the first storage controller, into the second storage device as well which is the object of data input/output effected by the second storage controller, the data frame having therein described information for specifying a Fiber Channel port connected to the first storage device or the second storage device which is the destination of the data frame and information for specifying the storage devices, wherein when the first storage controller receives the data frame in which is described a data writing request for the first storage device the first storage controller executes writing of data on that data frame to the first storage device while simultaneously the first storage controller creates a data frame in which the information for specifying the Fiber Channel port and the information for specifying the storage devices described in the data frame are rewritten such that the object is the second storage device, the first storage controller transmitting the created data frame to the second storage controller, and wherein the second storage controller receives the data frame and executes writing of data on that data frame to the second storage device, to thereby store a copy of the data stored in the first storage device into the second storage device as well that is the object of data input/output effected by the second storage controller.

According to this method, a mechanism for data copy management in which data stored in the first storage device is stored into the second storage device as well can be easily realized only by providing a mechanism for creating a data frame

as described above in the first storage controller.

It is to be noted that the function of the first storage controller relating to the data copy management is realized by, for example, a program that is run on the first storage controller.

[The Configuration of the Storage System]

The configuration of a storage apparatus system (storage system) described as an embodiment of the present invention is illustrated in Fig. 1. The storage system is operated in, for example, a data center. A host computer 5 and a first storage controller 10 are connected through first communication means 40 conforming to Fiber Channel Protocol. The host computer 5 is a computer such as, for example, a personal computer, a workstation and a general-purpose computer. The first storage controller 10 and a second storage controller 20 are connected through second communication means 50 conforming to the Fiber Channel Protocol.

The first storage controller 10 includes a switching control unit 11 and a storage device control unit 12. The switching control unit 11 and the storage device control unit 12 are connected by an internal communication path 13 such as an internal bus.

The switching control unit 11 comprises a switching device 111, a microprocessor 112, and a control memory 113 such as a ROM, a RAM, etc. The switching device 111 is provided with functions as a Fiber Channel switch (hereinafter referred to as "FC switch") and has a predetermined number of Fiber Channel ports. The switching device 111 receives a data frame transmitted from the host computer 5 and relays the data frame to the storage device control unit 12 or the second storage

controller 20. In addition, the switching device 111 also receives a data frame transmitted from the storage device control unit 12 or the second storage controller 20 and relays the data frame to the host computer 5. The microprocessor 112 provides
5 various functions of the switching control unit 11. In addition, the microprocessor 112 controls various components included in the switching control unit 11. Furthermore, the microprocessor 112 controls transmitting/receiving of the data frame conducted at the switching device 111. The control memory
10 113 is utilized as, for example, a work memory for processes executed by the microprocessor 112, or a storage area for various data such as tables and variables.

The storage device control unit 12 receives a data frame transmitted through the internal communication path 13 and
15 executes data input to and data output from a disk drive 14 in response to a data input/output request described in the data frame. The data input/output processes include a data writing request and a data read-out request. In the following description, data input and output includes the cases where
20 it includes either a data writing request or a data read-out request and where it includes both of a data writing request and a data read-out request.

The microprocessor 121 provides various functions of the storage device control unit 12, and controls the various
25 components included in the storage device control unit 12, etc. A control memory 123 is utilized as, for example, a work memory for the processes executed by the microprocessor 121 and as a storage area for various data such as tables and variables. A cache memory 124 is used as, for example, temporary storage
30 means for data being input to/output from the disk drive 14.

A drive controller 125 executes data input/output processes to the disk drive 14. The drive controller 125 may operate the disk drive 14 in RAID (Redundant Array of Inexpensive Disks) scheme. The drive controller 125 manages a physical storage area provided by the disk drive 14 by dividing it into logic volumes that are logic storage areas organized on the physical storage area. A specific identifier called LUN (Logical Unit Number) is given to each logic volume respectively. When an LUN has been designated, one or more disk drives 14 are identified. Furthermore, at least one or more logic volumes are related to one Fiber Channel port and, by designating one Fiber Channel port, all the logic volumes related to it are identified. The disk drive 14 can be held in the same housing as that of the first storage controller 10 or in housing different from that of the first storage controller 10.

The second storage controller 20 comprises a switching device 201, a microprocessor 202, a control memory 203, a cache memory 204, a drive controller 205 and a disk drive 21. The switching device 201 is provided with functions as a Fiber Channel switch (hereinafter referred to as "FC switch") having a plurality of Fiber Channel ports and second communication means 50 is connected to the Fiber Channel ports.

The microprocessor 202 provides various functions of the second storage controller 20 and controls various components included in the second storage controller 20, etc. The control memory 203 is utilized as, for example, a work memory for the processes executed by the microprocessor 202, or a storage area for various data such as tables and variables. The cache memory 204 is used as, for example, temporary storage means for data input to/output from the disk drive 21. The drive controller

205 executes data input/output processes to the disk drive 21.

The drive controller 205 may operate the disk drive 21 in RAID (Redundant Array of Inexpensive Disks) scheme. The drive controller 205 manages a physical storage area provided by the disk drive 21 by dividing it into logic volumes that are logic storage areas organized on the physical storage area. A specific identifier called LUN (Logical Unit Number) is given to each logic volume respectively. When an LUN has been designated, one or more disk drives 21 are identified.

10 The disk drive 21 can be held in the same housing as that of the second storage controller 20 or in housing different from that of the second storage controller 20.

[Data Frame]

15 Next, it will be described the structure of data frames conforming to the Fiber Channel Protocol, that is transmitted and received between the host computer 5, the first storage controller 10 and the second storage controller 20. The details of the Fiber Channel Protocol and data frames conforming to it are described in, for example, "Technical Instruction on Fiber Channel" edited by the Fiber Channel Technical Association, Ronso-sha, etc.

25 The structure of a data frame 70 on a Fiber Channel is illustrated in Fig. 2. At the head of the data frame 70, an SOF (Start of Frame) (71) of four bytes is placed and, following the SOF, a frame header (72) of 24 bytes, a data field (73) of 2112 bytes and a CRC (Cyclic Redundancy Check) (74) of four bytes are provided. At the tail of the data frame, an EOF (75) of four bytes is placed.

30 The structure of the frame header (72) of 24 bytes shown

in Fig. 2 is illustrates in Fig. 3. Destination_ID (81) is an identifier of the Fiber Channel port being the destination of the data frame and N_Port_ID, N_Port_Name or etc. is described in it. Source_ID (82) is an identifier of the Fiber Channel port of the source of the data frame and N_Port_ID, N_Port_Name or etc. is described in it.

The structure of FCP_CMND (Fiber Channel Protocol for SCSI Command) being one of the payloads of the data field shown in Fig. 2 is illustrated in Fig. 4. In FCP_LUN (FCP Logical Unit Number) (91), a LUN of the logic volume to which a SCSI command described in the payload is to be applied. In FCP_CNTL (FCP Control) (92), a command control parameter is described. In FCP_CDB (FCP Command Descriptor Block) (93), a SCSI command CDB (Command Data Block) is described. In FCP_DL (FCP Data Length) (94), the upper limit of the data size that can be converted between an initiator and a target is described.

[Basic Operation]

Next, the operation of the storage system while a data frame is transmitted from the host computer 5 to the first storage controller 10 will be described referring to a flow chart shown in Fig. 5.

First, a data frame is transmitted from the host computer 5 to a Fiber Channel port of the switching device 111 of the first storage controller 10 through the first communication means 40 (S511). When the switching device 111 of the first storage controller 10 receives the data frame (S512), the switching device 111 transmits an interrupt request to the microprocessor 112 (S513). When the microprocessor 112 receives the interrupt request, the microprocessor 112 stores

in the control memory 113 the contents of Destination_ID (81), Source_ID (82) and FCP_LUN (91) of the received data frame (S514).

5 In the control memory 113 of the first storage controller 10, a LUN management table shown in Fig. 6, in which whether the FCP_LUN (91) is directed to a logic volume of the first storage controller 10 or of the second storage controller 20 is described is stored. The microprocessor 112 checks whether the data frame is directed to a logic volume of the first storage controller or of the second storage controller by referring
10 the FCP_LUN (91) of the data frame stored in the control memory 113 to the LUN management table (S615).

Now, if the data frame is directed to a logic volume of the first storage controller 10, the first storage controller 10 transmits the data frame to the storage device control unit
15 12 through the internal communication path 13 (S516). When the storage device control unit 12 receives the data frame, the storage device control unit 12 executes data input and output corresponding to the SCSI command described in the data frame to the logic volume identified from the contents of FCP_LUN
20 (91) of the data frame (S517). After executing the data input/output process, the first storage controller 10 transmits to the switching control unit 11 the data frame in which the result of the process and data read out from the logic volume etc. are described, when necessary (S518). Then, when the
25 switching device 111 receives the data frame, it transmits the data frame to the host computer 5 (S519).

On the other hand, in (S515), if the data frame received from the host computer 5 is directed to a logic volume of the second storage controller 20, the microprocessor 112 controls
30 the switching device 111 such that the switching device 111

relays the data frame to the second storage controller 20 (S520).

When the second storage controller 20 receives the data frame

(S521), it executes data input and output corresponding to the

SCSI command described in the data frame to a logic volume

5 identified from the contents of FCP_LUN (91) of the data frame (S522).

After executing the data input/output process, the second storage controller 20 transmits to the first storage controller

10 through the second communication means the data frame in

10 which the report of completion, the result of the process or data read out from the logic volume are described, when necessary

(S523). In the Destination_ID (81) of the data frame,

information for specifying a Fiber Channel port of the host

computer 5 is described. When the switching device 111 of the

15 first storage controller 10 receives the data frame (S524), it relays the data frame to the host computer 5 (S525).

As described above, the data frame transmitted from the host computer 5 and input into the switching device 111 of the

first storage controller 10, is judged at the switching control

20 unit 11 whether it is directed to the first storage controller

10 or the second storage controller 20 and, if it is directed

to the second storage controller 20, the data frame is relayed

from the switching device 111 to the second storage controller

20. In other words, the first storage controller 10 receives

25 replacing the role of the second storage controller 20 the data

frame that the host computer 5 originally has directed to the

second storage controller 20.

Furthermore, by receiving the data frame by the first

storage controller 10 replacing the role of the second storage

30 controller 20, the host computer 5 only has to communicate with

the first storage controller 10 and the load of processing of the host computer 5 can be reduced by the amount to be created by the communication between the host computer 5 and the second storage controller 20.

5 Yet furthermore, by managing the LUN management tables unifying at the first storage controller 10, the logic volumes of the first storage controller 10 and the second storage controller 20 can be managed being unified and, therefore, the load of maintenance imposed on work such as updating of the
10 LUN management tables by operators etc. can be reduced. Especially when the disk drive 14 and disk drive 21 are operated while being virtualized, such a unified management scheme considerably contributes to more efficient operation and reduction of the load of operation.

15 When a disk array apparatus is the one equipped with a very ordinary interface as the connection interface with the host computer 5, like the switching device 201 as described above, the disk array apparatus can be used as the second storage controller 20 of the embodiment and, for example, in a storage
20 system using a disk array apparatus having a higher performance as the first storage controller 10 and a disk array apparatus having a lower performance than the other one as the second storage controller 20, a merit can be produced that the functions of the first storage controller 10 can be applied to the second
25 storage controller 20.

[Access Limit Function]

Next, the access function included in the first storage controller 10 will be described. The first storage controller
30 10 stores the access limit management table exemplified in Fig.

7. In the access limit management table, combinations of Destination_ID, Source_ID and LUN, that do not permit any data input/output process are described. The access limit function is a function that does not permit any data input/output process corresponding to a data input/output request described in the data frame when the combination of Destination_ID, Source_ID and LUN described in the data frame received by the first storage controller 10 from the host computer 5 is described in the access limit management table. The contents of the access limit management table may be registered automatically by a function of the first storage controller 10 or may be registered in a manual operation by an operator.

The access limit function will be described in detail referring to the flow chart shown in Fig. 8. This flow chart describes the case where a data frame in which a data input/output request directed to a logic volume of the second storage controller 20 is described is transmitted from the host computer 5 to the first storage controller 10.

When the first storage controller 10 receives a data frame (S812) transmitted from the host computer 5 (S811), the first storage controller 10 checks whether the combination of Destination_ID (81), Source_ID (82) and FCP_LUN (91) in the data frame is described in the access limit management table (S813, S814). If the combination is not described in the access limit management table, the first storage controller 10 relays the data frame to the second storage controller 20 (S815).

On the other hand, If the combination is described, the data frame is transmitted from the switching control unit 11 to the storage device control unit 12 (S816). When the storage device control unit 12 receives the data frame, the storage

device control unit 12 refers to the access limit management table in the control memory 113 (S817). If the combination of Destination_ID (81), Source_ID (82) and FCP_LUN (91) in the data frame is a combination not to be permitted the access, 5 the storage device control unit 12 creates a data frame directed to the host computer 5, in which it is described that it is not permitted to access, and transmits the newly created data frame to the switching device 111 (S818). Then, the newly created data frame is transmitted to the host computer 5 through 10 the switching device 111.

On the other hand, if the received data frame is a data frame to be permitted the access, the storage device control unit 12 transmits the data frame to the switching device 111 (S819) and the switching device 111 relays the data frame to 15 the second storage controller 20 (S820).

When the second storage controller 20 receives the data frame (S821), the second storage controller 20 executes data input and output corresponding to the SCSI command described in the FCP_CDB in the data frame to the logic volume described 20 in the FCP_LUN in the data frame (S822). After executing the data input/output process, the second storage controller 20 transmits to the first storage controller 10 the data frame in which the result of the process and data read out from the logic volume etc. are described, when necessary (S823).

25 In this embodiment, combinations of Source_ID, Destination_ID and LUN that do not permit any data input/output process are described in the access limit management table, but instead, combinations of Source_ID, Destination_ID and LUN that permit data input/output processes may be described in 30 the table. The access limit can be set in various forms including

not only the case where access limit is executed for all the combinations of Source_ID, Destination_ID and LUN but also the case where it is executed for any one of Source_ID, Destination_ID and LUN or for a combination of any two or more of these.

5 As described above, the first storage controller 10 executes replacing the role of the second storage controller 20 processes relating to the access limit. Therefore, it is enough that the function for executing the access limit is implemented on the first storage controller 10. Therefore, the
10 load of maintenance relating to the access limit function can be reduced by the amount to be created by providing such a mechanism to the second storage controller 20. Furthermore, because no access limit function needs to be included in the second storage controller 20, a storage controller having a
15 lower processing capacity than the first storage controller 10 can be used as the second storage controller 20 and, therefore, this is effective for reducing of the introduction cost and the operation cost.

 Furthermore, old types of storage controllers having lower
20 performance compared to the newly introduced storage controllers can be fully used as second storage controllers 20. Yet furthermore, it is enough that the access limit management table is stored only in the first storage controller 10 and, therefore, reduction of the load of maintenance created by the unified
25 management is progressed.

[Priority Control]

 Next, the priority control function included in the first storage controller 10 will be described. The first storage
30 controller 10 stores the priority management table shown in

Fig. 9. In the priority management table, the relations of Destination_ID, Source_ID and the delay time are described. The delay time is used as, for example, a parameter for determining the starting time of execution of a process that may increase the load of processing of the microprocessor 121. The priority control function is a function that controls the order of execution of the data input and output corresponding to the data frame received by the first storage controller 10 from the host computer 5, according to the delay time set in the priority management table. The contents of the priority management table may be registered automatically by a function of the first storage controller 10 or may be registered in a manual operation by an operator.

The priority control function will be described in detail referring to the flow chart shown in Fig. 10. When the first storage controller 10 receives a data frame A (S1012) transmitted from the host computer 5 (S1011), the first storage controller 10 checks whether the combination of Source_ID and Destination_ID in the data frame A exists in the priority management table or not (S1013, S1014).

If the combination exists, the first storage controller 10 determines the combination as the object for the access priority when a delay time is not set in the combination (S1015), executes data input and output for data frame A (S1016) and, after execution, transmits to the host computer 5 the data frame in which the report of completion is described (S1017).

If a delay time is set in (S1015), the first storage controller 10 checks whether another data frame B exists waiting to be processed in a queue managed by the control memory 123 in the storage device control unit 12 (S1020) and, if the data

frame B exists waiting, causes the starting of execution of the process for data frame A stand by for the period of the delay time set in the priority management table corresponding to the contents of the data frame A (S1021). The starting time
5 of which processes for the data frame should be delayed is adequately determined taking into consideration the load of processing imposed on the microprocessor 121 such that the data input/output process will be executed in an adequate order according to the needs of users. Arrangement may be such that
10 the control of the order of executing processes using the delay time as a parameter can be determined automatically by the storage system or that the control of the order can be set by a user from a control terminal connected to the storage system. After the delay time has passed, the first storage controller 10
15 executes data input and output for data frame A (S1022) and transmits to the host computer 5 a data frame in which the report of completion is described (S1023).

On the other hand, in (S1020), if another data frame B does not exist waiting in the queue, the first storage controller
20 10 executes processes according to the ordinary procedure for data frame A described in the basic operation above. That is, if the data frame A is directed to the disk drive 14 in the first storage controller 10, data input/output process for the disk drive 14 is executed in the first storage controller 10
25 and, if the data frame is directed to the disk drive 21 in the second storage controller 20, the first storage controller 10 transmits the data frame to the second storage controller 20, and the second storage controller 20 executes a data input/output process directed to the disk drive 21 (S1031). In either cases,
30 after the data input/output process has been completed, a report

of completion is transmitted if necessary (S1032).

In (S1014), if the combination of Source_ID, Destination_ID of the data frame A does not exist in the priority management table, the first storage controller 10 checks, according to
5 the procedure described in the above basic operation, whether the data frame A is directed to the disk drive 14 in the first storage controller 10 or to the second storage controller 20 (S1041). Now, if the data frame A is directed to the first storage controller 10, the data frame is transmitted to the storage
10 device control unit 12 (S1042) and the storage device control unit 12 executes a data input/output process for the data frame A (S1043) and transmits to the host computer 5 a data frame in which a report of completion is described (S1044).

On the other hand, if the data frame A is directed to the
15 disk drive 21 in the second storage controller 20, the first storage controller 10 transmits the data frame A to the second storage controller 20 (S1051). When the second storage controller 20 receives the data frame A (S1052), it executes a data input/output process directed to the disk drive 21 to
20 the data frame A (S1053) and transmits a data frame, in which a report of completion is described, to the host computer 5 through the first storage controller 10 (S1054).

As understood from the above description, the priority control of data frames in the host computer 5 is dedicatedly
25 executed in the first storage controller 10. That is, from another point of view, the first storage controller 10 executes replacing the role of the second storage controller 20 the priority control of the second storage controller 20. That is, it is enough that a priority control function is implemented
30 on the first storage controller 10, and the second storage

controller 20 can enjoy a function corresponding to the priority control function without having a mechanism for the function.

Furthermore, the load of operation can be reduced for the second storage controller 20 by the amount to be created by providing a mechanism for the priority control. Yet furthermore, it becomes possible to provide a function of priority control in the forms including the disk drive 21 in the second storage controller 20 even if a storage controller having a lower process capacity than that of the first storage controller 10 is used as the second storage controller 20 and, therefore, the introduction cost and the operation cost can be reduced.

Yet furthermore, old types of storage controllers having lower performance compared to the newly introduced storage controllers can be fully used as second storage controllers and, therefore, effective uses of resource and reduction of operation cost are progressed. In addition, it is enough that the priority management table is stored only in the first storage controller 10 and, therefore, reduction of the load of maintenance created by the unified management is progressed.

Though the priority control function for one LUN has been described as above, it can be arranged that a priority control is executed directing a plurality of LUNs and physical ports in a process described in one data frame.

[Data Copy Control]

In a storage system of the present invention, a copy management of data is conducted in a way in which a copy of the data stored in a logic volume (hereinafter referred to as "main logic volume") to which the first storage controller 10 executes data input/output processes is also stored in a logic

volume (hereinafter referred to as "sub-logic volume") in the second storage controller 20. The relation between the LUN in the main logic volume and the LUN in the sub-logic volume is stored in the control memory 113 in the first storage controller 10 as a pair management table. An example of the pair management table is shown in Fig. 11. However, the relation may be automatically set by a function of the first storage controller 10 or may be set in a manual operation by an operator etc.

Now, a process of the storage system while a data frame in which a data writing request directed to the main logic volume is described is transmitted from the host computer 5 to the first storage controller 10 will be described referring to a flow chart shown in Fig. 12.

When the first storage controller 10 receives the data frame (S1262) transmitted from the host computer 5 (S1261), it checks whether an LUN described in the data frame as FCP_LUN (91) exists in the pair management table (S1263, S1264).

If the LUN described in the data frame as FCP_LUN (91) does not exist in the pair management table, the first storage controller 10 transmits the received data frame to the storage device control unit 12 (S1265). When the storage device control unit 12 receives the data frame (S1266), it executes to the proper logic volume a data writing process corresponding to the SCSI command described in the data frame (S1267) and, when necessary, transmits to the host computer 5 a data frame in which a report of completion is described (S1268). Therefore, data has been written in the main logic volume.

In (S1264), if the LUN exists in the pair management table, the first storage controller 10 creates a new data frame in which the Destination_ID (81) and the FCP_LUN (91) of the data

frame is changed such that they designate the sub-logic volume being the destination of copying (S1269). The creation of the new data frame is executed in a way in which, after the received data frame has been stored in the control memory 113, the copy
5 of the data frame is stored in the control memory and the Destination_ID (81) and the FCP_LUN (91) in the copy of the data frame are changed to create a new data frame.

Next, the first storage controller 10 transmits the created new data frame to the second storage controller 20 by controlling
10 the switching device 111 to cause the switching device 111 to do so (S1270). When the second storage controller 20 receives the data frame (S1271), it executes a data writing process corresponding to the SCSI command described in the FCP_CDB of the data frame, to the sub-logic volume described in the FCP_LUN
15 of the data frame (S1272). In this way, data being a copy of the data described in the main logic volume has been also written in the sub-logic volume. After the data writing process has been executed, the second storage controller 20 transmits to the host computer 5 a data frame in which a report of completion
20 is described if necessary (S1273). In this way, data has been written in the main logic volume and, simultaneously, the same data has been written in the sub-logic volume.

On the other hand, in (S1264), if the LUN described in the data frame as the FCP_LUN (91) does not exist in the pair
25 management table, the first storage controller 10, following the procedure described in the above basic operation, checks whether the data frame is directed to the disk drive 14 in the first storage controller 10 or is directed to the second storage controller 20 (S1281). If the data frame is directed to the
30 first storage controller 10, the data frame is transmitted to

the storage device control unit 12 and the data writing process about the data frame is executed (1282). A data frame in which a report of completion is described is transmitted to the host computer 5 if necessary (S1284).

5 If the data frame is directed to the disk drive 21 in the second storage controller 20, the first storage controller 10 transmits the data frame to the second storage controller 20 (S1285). When the second storage controller 20 receives the data frame (S1286), it executes a data writing process directed
10 to the disk drive 21 (S1287). Then, the second storage controller 20 transmits to the host computer 5 through the first storage controller 10 a data frame in which a report of completion of the data writing process is described if necessary (S1288).

 In the above processes, when the first storage controller
15 10 receives a data frame in which a data writing request directed to the main logic volume is described, the first storage controller 10 creates a data frame for copying the request and transmits it to the second storage controller 20.

 According to this method, a mechanism of data copy
20 management, in which data stored in the disk drive 14 in the first storage controller 10 is also stored in the disk drive 21 in the second storage controller 20, can be easily realized only by providing a mechanism for creating a data frame as described above to the first storage controller 10.

25 There are two schemes for data copy management known such as so-called "synchronized scheme" in which the first storage controller 10 having received the data frame notifies the host computer 5 of a report of completion after confirming the completion of the writing processes for both of the main logic
30 volume and the sub-logic volume, and so-called "non-synchronized

scheme" in which the first storage controller 10 notifies the host computer 5 of the report of completion only when the writing process to the main logic volume has been completed regardless of whether the writing process to the sub-logic volume has been completed or not. However, the mechanism of the above-described embodiment can be applied to the case operated in either the "synchronized scheme" or the "non-synchronized scheme".

Although the storage system control method, etc., according to the present invention have been set forth by way of several embodiments, the above embodiments are only to facilitate the understanding of the present invention but are not intended to limit the present invention. It will be appreciated that the present invention can be variously changed or improved without departing from the spirit thereof and that the present invention covers equivalents thereof. Only one of the functions of access limit, priority, data copy control described above may be applied or any combination of two or more of these may be applied.

The communication protocol used for the first communication means or the second communication means is not limited to the Fiber Channel Protocol, but instead other protocols may be used.

In the above embodiments, a disk array apparatus has been described as an example of the storage apparatus. The present invention is however applicable to other storage apparatuses than the disk array apparatus, such as a semiconductor disk apparatus.

Although the present invention has been described above based on the embodiments, it is to be noted that the present invention shall not be limited to the embodiments stated herein, and that various modifications can be made without departing

from spirit of the invention.